



Contribution ID: 59

Type: **Student Talk**

The Geometry of Reinforcement Learning: Insights from the Dual Linear Program

Reinforcement Learning (RL) has become a cornerstone of machine learning, showcasing remarkable success in addressing real-world control problems and providing insights into cognitive processes in the brain. However, navigating the intricacies of modern RL proves challenging due to its numerous moving parts, escalating agent complexity, and the application of deep learning in a non-i.i.d. setting. The inherent challenge of intuitively reasoning about RL stems, in part, from its time-dependent and recursive nature. During this presentation, we explore the *dual linear program* and the intuitions it can offer. What traditionally serves as a theoretical construct for proving theorems emerges as a valuable tool for developing intuitions and facilitating the exploration of higher-level questions. We will focus on two practical demonstrations that underscore the significance of this perspective: 1) designing policy optimization algorithms and 2) pretraining RL agents. During the first half of this presentation, I will review the dual linear program and its geometry, aiming to uncover novel policy optimization strategies. In the second part, I will provide a preview of how the linear program can be generalized to *convex* MDPs, resulting in pretraining objectives similar to representation learning with the Variational Autoencoder.

Possible contributed talk

Yes

Are you a student?

Yes

Authors: MILOSEVIC, Nikola (Max Planck Institute for Human Cognitive and Brain Sciences); MÜLLER, Johannes (RWTH Aachen)

Co-authors: Dr SCHERF, Nico (MPI CBS); CAYÇI, Semih (RWTH Aachen)

Presenter: MILOSEVIC, Nikola (Max Planck Institute for Human Cognitive and Brain Sciences)

Session Classification: Student Session